

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
1 May 2003 (01.05.2003)

PCT

(10) International Publication Number  
**WO 03/036517 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**

(74) Agents: MEYER, Sheldon, R. et al.; Fliesler Dubb Meyer and Lovejoy LLP, Suite 400, Four Embarcadero Center, San Francisco, CA 94111-4156 (US).

(21) International Application Number: PCT/US02/33659

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(22) International Filing Date: 21 October 2002 (21.10.2002)

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(25) Filing Language: English

**Published:**

— with international search report

(26) Publication Language: English


(30) Priority Data:  
60/335,633 25 October 2001 (25.10.2001) US  
10/212,382 5 August 2002 (05.08.2002) US

(71) Applicant: BEA SYSTEMS, INC. [US/US]; 2315 North First Street, San Jose, CA 95131 (US).

[Continued on next page]

(72) Inventors: JACOBS, Dean, Bernard; 1747 Madera Street, Berkeley, CA 94707 (US). WOOLLEN, Rob; 2531 14th Avenue, San Francisco, CA 94127 (US). WHITE, Seth; 1045 Rivera Street, Apartment B, San Francisco, CA 94116 (US).

(54) Title: SYSTEM AND METHOD FOR FLUSHING BEAN CACHE



WO 03/036517 A1

(57) Abstract: Servers (110, 112, 113) in a network cluster (104) can each store a copy (116, 118, 120) of a data item (122) in local cache, providing read access to these copies through read-only entity bean. The original data item (122) in the database (108) can be updated through a read/write entity bean on one of the cluster servers. The cluster server has access to an invalidation target, which contains identification information relating to copies (116, 118, 120) of the data item stored on servers in the cluster. Once the read/write bean updates the data item in the database, an invalidate request can be sent or multicast to all cluster members, or to any read-only bean or server contained in the invalidation target. Each server or read-only bean receiving the request knows to drop any copy of the data item in local cache, and can request a current copy of the data item from the database.



*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**SYSTEM AND METHOD FOR FLUSHING BEAN CACHE****CLAIM OF PRIORITY**

[0001] This application claims priority to the following applications which are incorporated herein:

[0002] U.S. Provisional Patent Application No. 60/335,633, filed October 25, 2001, entitled "SYSTEM AND METHOD FOR FLUSHING BEAN CACHE," by Dean Bernard Jacobs and Rob Woollen.

[0003] U.S. Patent Application No. 10/212,382, filed August 5, 2002, entitled "SYSTEM AND METHOD FOR FLUSHING BEAN CACHE," by Dean Bernard Jacobs, Rob Woollen and Seth White.

10

**COPYRIGHT NOTICE**

[0004] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document of the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

**CROSS-REFERENCED CASES:**

[0005] The following applications are cross-referenced and incorporated herein by reference:

[0006] U.S. Provisional Application No. 60/305,986 entitled “DATA REPLICATION PROTOCOL,” by Dean Bernard Jacobs, Reto Kramer, and Ananthan Bala Srinivasan, filed July 16, 2001.

[0007] U.S. Provisional Application No. 60/316,187 entitled “CLUSTER CACHING WITH CONCURRENCY CHECKING,” by Dean Bernard Jacobs and Rob Woollen, filed August 30, 2001.

#### TECHNICAL FIELD

[0008] The invention relates generally to a system and method for storing data on a network.

#### BACKGROUND

[0009] When a data item is stored in a single database or data store that is accessible over a network, it is often the case that multiple servers or clients will require access to that data item. Traditionally, this requires a hit to the database each time the data item is accessed. Each hit to a database is relatively resource intensive and relatively inefficient.

[0010] One way of overcoming some of the efficiency and scalability problems is to store a local copy of the data item in cache memory. A server or client can then use that local copy if future access to the data item is needed. This process may be appropriate and efficient for data items that never change, but problems can arise when a data item is updated in the database.

[0011] If a data item in the database is updated, a copy of that data item stored in a local cache on the network will be different than the item in the database, as the cache will not automatically receive the update. The problem intensifies when there are local copies on multiple servers and/or clients on the

network. Since each of these local copies is created at a different time, there can be multiple versions of the data item on the network. If a user tries to update or view the data item, the copy accessed by the user may not be current and correct.

[0012] Such a problem with data latency can cause serious problems for applications that require near real-time accuracy, such as web sites that offer “real time” stock prices. Such an application might utilize a database table having at least two columns, one column containing stock symbols, which can be used as primary keys for the table, and one column containing the current price of each stock. In such an application, most of the activity involves users accessing the site and reading the current stock values. There is typically also activity involving back-end applications or systems that come in periodically, such as once every minute, with updated stock prices. These back-end systems need read/write access to the database in order to update the data.

[0013] Most access to the system will be read only. For these read-only users, the system can cache data to provide faster access. The system can update the cached information periodically, such as every fifteen minutes. In such a “read-mostly” situation, however, it may be preferable to give a user the most recent data. A fifteen minute delay in providing accurate information may be undesirable for many applications. It is typically desirable to give users information that is as accurate as possible.

[0014] One way to ensure that users get accurate information, or at least information that is current with data stored in the database, is to pull the information from the database for each request instead of reading a cached copy. This can be very expensive for many applications, as a hit to a database is much more time and resource intensive than reading a value from memory.

[0015] For people updating the data in the database, it may be desirable to wrap as many updates as possible into a batch transaction in order to improve performance. Wrapping updates into a single transaction also ensures that either

all the updates occur or none of the updates occur. Problems arise, however, in how to update cached copies for each item updated in a transaction.

#### BRIEF SUMMARY

**[0016]** A system and method are included for updating a copy of a data item stored in local cache on at least one server in a network cluster. Identification information is provided to a read/write bean stored on a server in the cluster. The identification information relates to any server in the cluster that contains a read-only bean and a copy of the data item in local cache. A read-only bean provides read access to the local copy of the data item. The original data item is stored in a network database, and is updated using the read/write bean. When the data item is updated by the read/write bean, an invalidate request can be sent or multicast from the server containing the read/write bean to the entire cluster, or can be sent to any server or read-only bean identified by the identification information having a local copy of the data item. Any local copy of the data item can then be dropped in response to the request. A current copy of the data item can be read from the database and stored in local cache.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0017]** Figure 1 is a diagram of a system in accordance with one embodiment in accordance with the present invention.

**[0018]** Figure 2 is a diagram of a system in accordance with the embodiment of Figure 1.

**[0019]** Figure 3 is a diagram of an alternative embodiment of the system of Figure 1.

[0020] Figure 4 is a flowchart for a method in accordance with one embodiment of the present invention.

[0021] Figure 5 is a flowchart for a method in accordance with another embodiment of the present invention.

[0022] Figure 6 is a diagram of a system in accordance with another embodiment in accordance with the present invention.

#### DETAILED DESCRIPTION

[0023] In order to maintain consistency among items distributed on a network, a system in accordance with the present invention can take advantage of beans, or JavaBeans. A bean is basically a framework for components that can be added to a server to extend functionality. One embodiment utilizes two types of beans, “read-only” entity beans and “read/write” entity beans. An entity bean is a bean that is persistent, allows shared access, has primary keys, and can participate in relationships with other entity beans. Each entity bean can have an underlying table in a relational database, with each instance of the bean corresponding to a row in that table.

[0024] A read-only bean is a bean that can be cached on a server, such as an enterprise JavaBean that resides in a network cluster. The read-only bean can provide read access to any server in the cluster, as well as to any client inside or outside of the cluster. The read/write bean is transactional, residing on a server in the cluster and providing cluster servers with read/write access to a network database. The read-only bean deals with data in local cache on a cluster server. The read/write bean deals with information in a database.

[0025] One way to address the concurrency of information in the cache and in the database is to associate a timeout value with each read-only entity bean. For example, a read-only bean can be deployed with a default cycle of ten minutes.

After each period of ten minutes passes, the read-only bean goes back to the database and reads the current value. This approach can work well for certain applications, such as those applications with values that change at a regular interval.

[0026] There may be applications, however, which have data that changes very infrequently. When this data changes, users may want to know about the change as soon as possible. Since the data does not change very often, it is tempting to set a long read cycle time in order to conserve resources. This can have the undesirable effect, however, of creating latency issues with the data, as the resultant delay in updating the data can be almost as long as the cycle time, depending on the point in the cycle at which the update occurs. For such applications, it is desirable that the data accessible by a read-only user is updated as soon as possible after the data in the database is updated.

[0027] One system in accordance with the present invention provides an interface, exposed by a read-only bean. The interface allows a user or application to tell the system to drop a cache, or “invalidate” a cache, when the user updates a data item or is aware of an update. This interface shall be referred to as a “CachingHome,” as an entity bean typically has a “home” or factory that creates it. CachingHome can have three methods on it, and be coded as follows:

```
package weblogic.ejb;
public interface CachingHome {
    public void invalidate(Object pk) throws RemoteException;
    public void invalidate(Collection pks) throws RemoteException;
    public void invalidateAll() throws RemoteException;
}
```

[0028] The method invalidate(Object pk) lets a user invalidate data associated with a particular primary key in a database or data table. The method invalidate(Collection pks) lets a user invalidate data for a collection or group of keys. The method invalidateAll() allows a user to invalidate data for all keys in the database table. These invalidate methods allow a user to ensure that values are stored in local cache until a programmer, application, user, or system says otherwise.

[0029] One such method 300 is shown in the flowchart of **Figure 5**. A copy of a data item is stored on at least one server in a network cluster 302. The data item can be updated in the database using a read/write bean on one of the cluster servers 304. An invalidate request can then be initiated using an interface of a read-only bean located on one of the servers, the request being sent to any server containing a local copy of the data item 306. Any copy of the data item can then be dropped from a server receiving the request 308.

[0030] In a system 100 with a network cluster 104, such as is shown in **Figure 1**, it is possible that a copy of a value 122 stored in a database 108 is cached on each server 110, 112, 114 in the cluster 104. If a client 102 contacts server 110 through the network 106 and requests that server 110 invalidate a given key, such as by making the request “invalidate (Key),” it is easy for server 110 to drop its cached copy 116 of the value 122 or values associated with that key. A problem exists, however, in how to get servers 112 and 114 to drop their cached copies 118, 120 as well.

[0031] One embodiment allows server 110 to drop a copy 116 in local cache when it receives an invalidate request 124 from the client 102, as shown in **Figure 2**. After dropping the copy 116 from local memory, server 110 can send a message 126 on multicast to the other servers 118, 120 or read-only beans in the cluster 104 to drop the copy of the value in local cache. Multicast is a technique

for sending a packet of information or a message from one server, or source, to other servers without having any unnecessary packet duplication, as one packet is sent from the source and replicated across the network as needed. This approach allows each server to drop the cached value when the value in the database is updated.

[0032] **Figure 6** shows another system which includes multiple clients **402** and multiple databases **416** in communication with the network **404**. A data item **418** is stored in one of the databases **416**. A copy **420** of the data item is stored in a cluster server **406** of cluster **408**. There are two copies **422** stored in cluster servers **422** of cluster **412**, as well as a copy **424** on server **414**, which is not contained in any server cluster. This system would work similar to the system of Figures 1 and 2, in that one of the servers **406**, **410**, **414** containing a copy of the data item **418** can drop a copy in local cache when it receives an invalidate request from a client, and can send a message on multicast to the other servers on the network to drop any copy of the value in local cache.

[0033] Another problem exists due to the fact that a multicast message is only sent once by the source and does not wait for confirmation of receipt by the other servers. A server in the cluster might not get an invalidate request, such as if it is temporarily offline. A system in accordance with the present invention can provide a more reliable multicast by tagging each such message or request with a version number or sequential number. In this way, a server receiving a request will know the version of the request, as well as the version of the last request it received, such that the server will know if it missed a message. Once a server determines that it has missed a message, it can request that the message be sent again so that it can update accordingly.

[0034] One problem with this approach, however, is that a server will not know it has missed an update until another update is sent. In certain applications

such as an on-line store posting weekly specials for weeks 1, 2, and 3, it may be unacceptable to wait until the next update to get correct information. The store would not realize that it had missed the specials for week 2 until the update for week 3. The week 1 specials would have remained up during week 2, displaying the wrong information to any user accessing the system during that time. When the system realizes that it missed the week 2 update, it will already be week 3. The server will end up simply discarding the week 2 information without the information ever having been displayed to a user.

**[0035]** A system in accordance with the present invention can get around this problem by periodically “heartbeating” information to the servers in the cluster. A server heartbeats a packet of information or a message by sending the message periodically over the network or cluster. A heartbeat message can contain information such as the latest version number, previous version numbers, or the actual update information itself if the update information is small enough to be practical. If a server receives a heartbeat message containing the latest version number, and the server is not on that version of the data or did not receive the latest invalidate request, the server can request or pull the invalidate message from the server.

**[0036]** The initiating server that initially sent the invalidate request, which may also be the server sending the multicast and/or heartbeats, can store recent requests for a certain amount of time or can store a certain number of recent requests. If a cluster server requests an invalidate message that the initiating server is still storing, the initiating server can simply send the message to the cluster server, by a method such as a multicast or a point-to-point connection. If the initiating server no longer has the message, the initiating server can tell the cluster server to simply drop its entire cache, since it is not possible to tell the cluster server which keys have changed. The cluster server can read new and/or current information from the database. This can temporarily lessen performance, but the

newly-cached information will at least be current with the information in the database.

[0037] In operation, a client or application can update a data item through a read/write entity bean. The update, or a transaction containing multiple updates, will commit to the database. An invalidate message can be sent to the servers in the cluster, with the message being triggered for example by the client or server updating the data item. The cluster servers can each drop any copy in local cache and can read in the new value from the database, either right away, later, or when necessary to serve a request. Normally it is not possible to read uncommitted data, so it may be preferable to use a two-step process where the data is committed first and then a message is multicast to the cluster.

[0038] One problem with the above approach is that it forces a client to initiate an invalidate request, which can involve a little more complexity for the client. There is also the possibility that the client could use the invalidate method incorrectly or make a mistake. It may therefore be preferable that the system can do it automatically.

[0039] A system in accordance with the present invention can address this problem by using an “invalidation target.” An invalidation target is based on the idea that the read-only and read/write beans point to the same data in the database, with the people reading the data using the read-only bean and the people updating the data using the read/write bean. The idea is to invalidate the read-only bean when the read/write bean is updated or modified.

[0040] When deploying an entity bean or enterprise JavaBean, there is typically a deployment descriptor used to store meta data about the actual entity bean. A deployment descriptor can be, for example, an XML document used to provide information about the services provided by the bean that should be made available to clients. The information can provide a wide variety of information to the clients, such as request routing information, as well as method and class details

for a supporting Java class. A tag can be added to the deployment descriptor, referred to previously as an “invalidation target.” The invalidation target for a read/write bean can contain the identity of any associated read-only bean. The invalidation target can be used to automatically invalidate any associated read-only bean(s) when the read-only bean is updated.

[0041] In one embodiment, the invalidation target can be updated when a server requests information from a read/write bean or generates a read-only bean does the request. When the read/write bean forwards information from a database or data store to the requesting server, the read/write bean can also update the invalidation target. An XML file stored on the server containing the read/write bean can be updated to include the identity of the server requesting the information or creating the bean.

[0042] **Figure 3** again shows the system 100 of Figure 1, except in this embodiment the system is shown taking advantage of an invalidation target. Whenever a read/write bean 128 is used to update a data item 122, the system can look to the invalidation target associated with that read/write bean 128 and send an invalidate request to each read-only bean 130 associated with that read/write bean. The invalidate request can be acted on directly by server 110, which contains both the read/write bean 128 and a read-only bean 130. The invalidate request can also be sent by any appropriate protocol to other servers containing a read-only bean 130, which would be in the invalidation target. In one approach, server 110 contacts server 112 directly by a point-to-point connection 134, telling the read-only bean 130 on server 112 to drop the cached copy 118 on server 112. In another approach, server 110 can send a multicast message 132 over the network 106 to any server which contains a read-only bean 130 within the scope of the invalidation target, such as server 112.

[0043] A subsequent call to an invalidated read-only bean can cause a method such as ejbLoad to be called, which can read current information from the database into cache. For example, a container-managed persistence (CMP) bean, an entity bean whose state is automatically synchronized with a database, can use an invalidation-target element in an XML file such as ejb-jar.xml to specify a read-only entity bean that should be invalidated when the read/write bean has been modified. The container in this example can automatically invalidate the invalidation-target, such as after the transaction is completed.

[0044] Such a method 200 is shown in the flowchart of **Figure 4**. A copy of a data item is stored on at least one server in a network cluster 202. Identification information, or an invalidation target, is provided to a read/write bean on a server in the cluster, the information relating to any server in the cluster containing a read-only bean and a local copy of the data item 204. The data item can be updated in the database using a read/write bean 206. An invalidation request can be sent from the server containing the read/write bean to any server identified by the identification information 208. The copy of the data item can then be dropped by each server receiving the request 210.

[0045] In this way, customers or clients do not have to write any additional code to invalidate an item. In accordance with an embodiment of the present invention, only an invalidation target must be specified in order to keep the read/write and read-only beans coherent. The beans can coexist on the same server, with the read-only bean reading items from local cache and the read/write bean reading from, and writing to, the database.

[0046] In order to improve performance, a system in accordance with the present invention can instead wait until an entire transaction or series of updates is committed to the database or data table, instead of sending an individual message for each update. A server, such as the server initiating the update, can keep track

of which keys were updated during the transaction and multicast a single message that includes information for all updated primary keys. Such batching of messages can improve the overall performance of such a system and can reduce the opportunity for error or inconsistencies.

**[0047]** One example of a system that can be used in accordance with the present invention contains a table of stock symbols, as well as information associated with each symbol, such as price and volume. A Java server page can be used to allow a user to request the current price of a stock. The Java server page can read the information from a read-only entity enterprise Java bean. A Java Message Service (JMS) queue can receive messages with updates to stock prices. An message-driven bean can de-queue these messages and update the associated CMP entity bean. When this modification occurs, the container can invalidate the associated read-only bean.

**[0048]** The foregoing description of preferred embodiments of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to one of ordinary skill in the relevant arts. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

CLAIMS

What is claimed is:

1        1. A method for updating a cached copy of a data item in a database,  
2 comprising:

3                 storing a copy of a data item on a server, the data item being stored in a  
4 database;

5                 generating a read-only bean to provide read access to the copy and storing  
6 identification information relating to the read-only bean;

7                 generating a read/write bean to provide read and write access to the data  
8 item in the database and updating the data item using the read/write bean; and

9                 sending an invalidate request to the read-only bean using the identification  
10 information so the copy of the data item can be deleted.

1        2. A method according to claim 1, wherein:

2                 storing a copy of a data item stores the copy in cache on a first server.

1        3. A method according to claim 2, wherein:

2                 generating a read-only bean generates the read-only bean on the first server;

3                 and

4                 generating a read/write bean generates the read/write bean on a second  
5 server.

1        4. A method according to claim 3, wherein:

2                 sending an invalidate request sends the request from the second server  
3 containing the read/write bean to the first server containing the read-only bean.

1        5. A method according to claim 1, wherein:

2           generating a read-only bean and storing identification information stores the  
3       identification information as a tag in an XML document.

1       6.     A method according to claim 1, wherein:  
2           sending an invalidate request to the read-only bean further involves  
3       multicasting the request to any read-only bean storing a copy of the data item as  
4       identified by the identification information.

1       7.     A method according to claim 1, wherein:  
2           sending an invalidate request to the read-only bean further involves  
3       multicasting the request to any member of a cluster, the server being a member of  
4       the cluster.

1       8.     A method according to claim 7, further comprising:  
2           periodically multicasting the invalidate request across the cluster.

1       9.     A method according to claim 1, wherein:  
2           sending an invalidate request to the read-only bean occurs by a point-to-  
3       point connection.

1       10.    A method according to claim 1, wherein:  
2           sending an invalidate request involves multicasting a request containing a  
3       version number.

1       11.    A method according to claim 1, further comprising:  
2           including a version identifier in the invalidate request.

1       12.    A method according to claim 11, further comprising:

2           determining whether an invalidate request for the read-only bean was  
3       missed by checking the version identifier; and  
4           requesting that an invalidate message be resent to the read-only bean if the  
5       read-only bean missed an invalidate request.

1     13.   A method according to claim 1, further comprising:  
2           storing recent invalidate requests.

1     14.   A method according to claim 1, further comprising:  
2           collecting information for multiple updates and batching the update  
3       information into one invalidate request.

1     15.   A method for updating a cached copy of a data item, comprising:  
2           storing a copy of a data item on a server in a cluster, the data item being  
3       stored in a network database;  
4           updating the data item in the network database;  
5           sending an invalidate request to a read-only bean, the read-only bean  
6       providing the cluster with read access to the copy; and  
7           deleting the copy of the data item from the server in response to the  
8       invalidate request.

1     16.   A method for updating a cached copy of a data item, wherein:  
2           storing a copy of a data item stores the copy on a first server in the cluster,  
3       the read-only bean being stored on a second server.

1     17.   A method for updating a cached copy of a data item, comprising:  
2           storing a copy of a data item on a server in a cluster, the data item being  
3       stored in a cluster database and having a state;

4           providing the cluster with read access to the copy through a read-only bean  
5       on the server;  
6           setting a timeout period for the read-only bean; and  
7           reading the state of the data item in the cluster database at the end of the  
8       timeout period and updating the copy of the data item.

1       18.   A system for updating cached copies of a data item in a cluster, comprising:  
2           a first cached copy of the data item on a first server in the cluster;  
3           a first read-only bean allowing read access to the first cached copy;  
4           a second cached copy of the data item on a second server in the cluster;  
5           a second read-only bean allowing read access to the second cached copy;  
6       and  
7           a read/write bean allowing write access to the data item, the read/write bean  
8       programmed to send an invalidation request to the first and second read-only beans  
9       when the read/write bean updates the data item.

1       19.   A system according to claim 18, further comprising:  
2           a database adapted to contain the data item.

1       20.   A system according to claim 18, wherein:  
2           the first read-only bean has an interface allowing the first cached copy of  
3       the data item to be deleted; and  
4           the second read-only bean has an interface allowing the second cached copy  
5       of the data item to be deleted.

1       21.   A system according to claim 18, wherein:  
2           each of the first and second read-only beans is adapted to request the  
3       resending of the invalidation request if the request was not received.

- 1        22.     A system according to claim 21, wherein:  
2                  each of the first and second read-only beans is adapted to read the data item  
3                  if unable to obtain a resending of the invalidation request.
  
- 1        23.     A system according to claim 18, further comprising:  
2                  an invalidation target containing the identity of each read-only bean in the  
3                  cluster caching a copy of the data item.
  
- 1        24.     A system according to claim 23, wherein:  
2                  the invalidation target further contains the identity of the read/write bean.
  
- 1        25.     A system according to claim 23, wherein:  
2                  the read/write bean is adapted to examine the invalidation target in order  
3                  to send an invalidation request to each read-only bean caching a copy of the data  
4                  item.
  
- 1        26.     A system according to claim 18, wherein:  
2                  each of the first and second read only beans is further adapted to read a new  
3                  copy of the data item after processing an invalidation request.
  
- 1        27.     A system for updating cached copies of a data item, comprising:  
2                  a first server containing a first cached copy of the data item, the first server  
3                  containing a first read-only bean allowing read access to the first cached copy; and  
4                  a second server containing a second cached copy of the data item, the second  
5                  server containing a read/write bean and a second read-only bean, the second  
6                  read-only bean allowing read access to the second cached copy, the read/write bean  
7                  allowing write access to the data item, the read/write bean adapted to send an

8        invalidation request to the first and second read-only beans when the read/write  
9        bean updates the data item in the data storage medium.

1        28.      A system according to claim 27, wherein:  
2                  the first read-only bean is adapted to delete the first cached copy of the data  
3                  item and the second read-only bean is adapted to delete the second cached copy of  
4                  the data item in response to the invalidation request.

1        29.      A system for updating cached copies of a data item, comprising:  
2                  a first read-only bean allowing read access to a first cached copy of the data  
3                  item;  
4                  a second read-only bean allowing read access to a second cached copy of  
5                  the data item; and  
6                  a read/write bean allowing write access to the data item, the read/write bean  
7                  adapted to send an invalidation request to the first and second read-only beans  
8                  when the read/write bean updates the data item.

1        30.      A system for updating cached copies of a data item, comprising:  
2                  a first server containing a first cached copy of the data item, the first server  
3                  containing a read/write bean and a first read-only bean allowing read access to the  
4                  first cached copy, the read/write bean allowing write access to the data item and  
5                  adapted to send an invalidate request to the first read-only bean when the read/write  
6                  bean updates the data item; and  
7                  a second server containing a second cached copy of the data item, the  
8                  second server containing a second read-only bean allowing read access to the  
9                  second cached copy, the first read-only bean capable of sending a request to the  
10                 second server to delete the second cached copy.

- 1        31.     A system for updating cached copies of a data item, comprising:
  - 2              a first server containing a first cached copy of the data item, the first server
  - 3              containing a read/write bean and a first read-only bean, the read/write bean
  - 4              allowing write access to the data item in the data storage medium, the first read-
  - 5              only bean allowing read access to the first cached copy, the first read-only bean
  - 6              containing a timeout value specifying how often the first server update the cached
  - 7              copy of the data item; and
  - 8              a second server containing a second cached copy of the data item, the
  - 9              second server containing a second read-only bean allowing read access to the
  - 10             second cached copy, the second read-only bean containing a timeout value
  - 11             specifying how often the first server update the cached copy of the data item.
- 1        32.     A system for updating cached copies of a data item in a network cluster,  
2           comprising:
  - 3              a database table containing a data item;
  - 4              a client adapted to make requests to the network cluster;
  - 5              a first server in the network cluster containing a first cached copy of the
  - 6              data item, the first server containing a read-only entity bean allowing read access
  - 7              to the first cached copy; and
  - 8              a second server in the network cluster containing a second cached copy of
  - 9              the data item, the second server containing a read/write entity bean and a second
  - 10             read-only entity bean, the second read-only entity bean allowing read access to the
  - 11             second cached copy, the read/write entity bean allowing write access to the data
  - 12             item in the database table and containing identification information for the first and
  - 13             second read-only entity beans, the read/write entity bean adapted to send an
  - 14             invalidation request to the first and second read-only entity beans when the
  - 15             read/write entity bean updates the data item in the database table in response to a
  - 16             request from the client.

1       33. A computer-readable medium, comprising:  
2           means for storing a copy of a data item on a server, the data item being  
3           stored in a database;  
4           means for generating a read-only bean to provide read access to the copy  
5           and storing identification information relating to the read-only bean;  
6           means for generating a read/write bean to provide read and write access to  
7           the data item in the database and updating the data item using the read/write bean;  
8           and  
9           means for sending an invalidate request to the read-only bean using the  
10          identification information so the copy of the data item can be deleted.

1       34. A computer program product for execution by a server computer for cached  
2           copies of a data item in a network cluster, comprising:  
3           computer code that can store a copy of a data item on a server, the data item  
4           being stored in a database;  
5           computer code that can generate a read-only bean to provide read access to  
6           the copy and storing identification information relating to the read-only bean;  
7           computer code that can generate a read/write bean to provide read and write  
8           access to the data item in the database and updating the data item using the  
9           read/write bean; and  
10          computer code that can send an invalidate request to the read-only bean  
11          using the identification information so the copy of the data item can be deleted.

1       35. A system for updating a data item on a network, comprising:  
2           means for storing a copy of a data item on a server, the data item being  
3           stored in a database;

4           means for generating a read-only bean to provide read access to the copy  
5       and storing identification information relating to the read-only bean;

6           means for generating a read/write bean to provide read and write access to  
7       the data item in the database and updating the data item using the read/write bean;  
8       and

9           means for sending an invalidate request to the read-only bean using the  
10      identification information so the copy of the data item can be deleted.

1     36.   A computer system comprising:

2           a processor;

3           object code executed by said processor, said object code configured to:

4               store a copy of a data item on a server, the data item being stored in  
5       a database;

6               generate a read-only bean to provide read access to the copy and  
7       storing identification information relating to the read-only bean;

8               generate a read/write bean to provide read and write access to the  
9       data item in the database and updating the data item using the read/write bean; and  
10               send an invalidate request to the read-only bean using the  
11      identification information so the copy of the data item can be deleted.

1     37.   A system for updating a data item in a server cluster, comprising:

2           a network client capable of making a request over the network;

3           a network database for storing a data item and providing access to that data  
4       item over the network; and

5           a network server for receiving the request from said network client and  
6       processing the request, said network server storing a local copy of the data item to  
7       be used in processing the request, the network server containing a read/write bean  
8       and a read-only bean, the read-only bean allowing read access to the local copy, the

9       read/write bean allowing write access to the data item in the network database and  
10      containing identification information for the read-only bean, the read/write bean  
11      adapted to send an invalidation request to the read-only bean on the network server  
12      and any read-only bean on any other server in the server cluster when the read/write  
13      bean updates the data item in the network database.

1       38.     A system for updating a data item in a server cluster, comprising:  
2               a plurality of network clients, each client being capable of making a request  
3               over the network;  
4               at least one network database for storing a data item and providing access  
5               to that data item over the network; and  
6               a plurality of network servers, each network server being capable of  
7               receiving a request from said plurality of network clients and processing the  
8               request, each network server capable of storing a copy of the data item in local  
9               cache and capable of storing a read-only bean for providing access to the copy of  
10          the data item to be used in processing a request from one of the network clients,  
11          one of said plurality of network servers containing a read/write bean allowing  
12          write access to the data item in the network database and having access to  
13          identification information for each read-only bean, the read/write bean being  
14          programmed to send an invalidation request to each read-only bean when the  
15          read/write bean updates the data item in the network database such that each read-  
16          only bean can delete a copy of the data item in local cache.




Figure 1

2/6




Figure 2

3/6






Figure 3



*Figure 4*

300

5/6



*Figure 5*




Figure 6

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US02/33659

**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(7) : G06F 17/30  
US CL : 707/8, 200, 201

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
U.S. : 707/8, 200, 201

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
WEST

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,999,947 A (ZOLLINGER et al) 07 December 1999 (07.12.1999), figure 5, column 3 lines 31-67	1-38
Y	US 6,298,478 B1 (NALLY et al) 02 October 2001 (02.10.2001), note figures 5 and 6A	1-38
Y,P	US 6,449,734 B1 (SHRIVASTAVA et al) 10 September 2002 (10.09.2002), abstract, figures 3 and 4, column 2 lines 8-46	1-38
Y	US 5,897,634 A (ATTALURI et al) 27 April 1999 (27.04.1999), column 10 lines 7-67	1-38
Y	US 5,546,579 A (JOSTEN et al) 13 August 1996 (13.08.1996), abstract, column 9 line 62 through column 11 line 63.	1-38

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"E" earlier application or patent published on or after the international filing date

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"O" document referring to an oral disclosure, use, exhibition or other means

"&"

document member of the same patent family

"P" document published prior to the international filing date but later than the priority date claimed

Date of the actual completion of the international search

07 December 2002 (07.12.2002)

Date of mailing of the international search report

20 DEC 2002

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks  
Box PCT  
Washington, D.C. 20231

Authorized officer

Greta L. Robinson

Facsimile No. (703)305-3230

Telephone No. (703) 305-9600